

ABSTRACT OF THE DISCLOSURE

In a distributed file system the distributed storage management is made useful to a variety of applications. Multiple quality of service options are provided through locking. Three locking systems are provided. The system offers a locking system designed for sequential consistency with write-back caching, typical of distributed file systems. A second locking system is provided for sequential consistency with no caching for applications that manage their own caches. Finally, a locking system that implements a weaker consistency model with write-back caching, designed for efficient replication and distribution of data is included. Locks for replication are suitable for serving dynamic data on the Internet and other highly-concurrent applications. The selection of the appropriate lock protocol for each file is set using the file metadata. Further, a novel locking system is provided for the lock system implementing a weak consistency model with write back caching. This system is implemented utilizing two whole file locks: a producer lock P and a consumer lock C. Any client can hold a consumer lock and when holding a consumer lock can read data and cache data for read. The producer lock is only held by a single writer and a writer holding a producer lock can write data, allocate and cache data for writing. When a writer performs a write, the write is performed as an out-of-place write. An out-of-place write writes the data to a different physical storage location than from which it was read. By performing an out-of-place write the old data still exists and is available to clients. Once the writer completes the write and releases the producer lock the previous data is invalidated and the clients are informed of the new location of the data. Clients can then read the new data from storage when needed and the server reclaims the old data blocks.